

MEMORY MANAGEMENT FOR SYSTEMS FOR GENERATING  
3-DIMENSIONAL COMPUTER IMAGES

This invention relates to memory management for use with systems for generating 3-dimensional computer generated images.

In our British Patent No. 2281682, there is described a 3-D rendering system for polygons in which each object is seen to be viewed as defined in a set of surfaces which are infinite. Each elementary area of the screen in which an image is to be displayed has a ray projected through it from a viewpoint into the 3-dimensional scene. The location of the intersection of the projected ray with each surface is then determined. From these intersections, it is then possible to determine whether any intersected surface is visible at that elementary area. The elementary area is then shaded for display in dependence on the results of the determination.

The system can be implemented in a pipeline type processor comprising a number of cells, each of which can perform an intersection calculation with a surface. Thus, a large number of surface intersections can be computed simultaneously. Each cell is loaded with a set of coefficients defining the surface for which it is to perform the intersection test.

A further improvement is described in our UK Patent Application No. 2298111. In this, the image plane is subdivided into sub-regions or tiles such that tiles can be processed in turn. It is proposed to use a variable tile size and to project a bounding box around complex objects so that only those tiles falling within the bounding box require processing. This is done by firstly determining the distribution of objects on the visible screen in order for a suitable tile size to be selected. The surfaces which define the various objects are then stored in a list, known as the display list, thereby

avoiding the need to store identical surfaces for each tile, since one object made of many surfaces could appear in a number of tiles. Object pointers which identify the objects in the display list are also stored. There is one  
5 object pointer list per tile. The tiles can then be rendered in turn using the ray casting technique described above until all objects within each tile are processed. This is a useful method because no effort needs to be made to render objects which are known not to be visible in a  
10 particular tile.

A further improvement on this is proposed in our International Patent Application No. PCT/GB99/03707, in which any tiles within the bounding box which are not required to display a particular object are discarded  
15 before rendering.

A block diagram of the type of processor used is shown in Figure 1. This comprises a tile accelerator unit 2 which performs the tiling operations referred to above and supplies the tile information to a display list memory  
20 4. This in turn supplies tile object data for processing to the image synthesis processor (ISP) 6 which performs the ray/surface intersection tests referred to above. This derives image data which comprises object identification and depth data. After this, the thus derived image data  
25 is supplied to a texturing and shading processor (TSP) 8 which applies texturing and shading data to surfaces which have been determined as visible and outputs image shading data to a frame store.

Most 3-D graphic systems use what is known as  
30 immediate mode rendering. This type of rendering avoids the needs to store a copy of the entire scene in the display list memory. Instead, each polygon is rendered the moment that it is presented to the rendering hardware. This has the advantage of avoiding the storage  
35 requirements of the display list memory. However, it also precludes the considerable optimisation which are possible

with the type of rendering described in our British Patent No. 2281682 and improved on in our British Patent Application No. 2298111. Advantages such as deferred texturing and translucency sorting are possible using this type of architecture.

The complexity of scenes which it is necessary to render using 3-D graphic systems is increasing as hardware becomes more powerful. However, this presents a problem for display list based renderers because the display list storage requirements also increase. In the past, if the display list memory became exhausted, parts of the scene would simply not be rendered. Image corruption of this sort is no longer considered acceptable. The cost of memory, which is not falling as fast as scene complexity is rising, makes increasing the amount of memory in the system unfeasible.

Preferred embodiments of the present invention provide a method of managing the display list in such a manner as to impose an upper bound on the memory consumption, whilst at the same time attempting to minimise the memory bandwidth consumed by the system. This may be achieved by storing the state of the system (ISP and TSP) to memory before rendering of a tile is complete, and reloading this state at a later time in order to finish the rendering. We refer to this technique as "z/frame buffer load and store".

In a preferred embodiment, this screen is divided up into a number of regions called macro-tiles, in which each of these consists of a rectangular region of the screen composed of a number of smaller tiles. Memory in the display list is then divided into blocks and these are listed in a free store list. Blocks from the free store are then allocated to the macro-tiles as required. The tiling operation stores polygon parameter data and object pointers for surfaces in each block associated with each macro-tile in which they are visible. When the memory for

the parameters fills up, or reaches some predefined threshold, the system selects a macro-tile, performs a z/frame buffer load, and renders the contents of the macro-tile before saving it using a z/frame buffer store operation. Upon completion of such a render, the system frees any memory blocks associated with that macro-tile, thereby making them available for further parameter storage. The z/frame buffer load and store operations are restricted to the macro-tiles that are actually rendered rather than every tile on the screen as was previously the case. Tiling of either the remainder of the current frame or of the next frame to be displayed then continues in parallel with macro-tile renders and the allocation of blocks of memory from the same pool to further macro-tiles. It is not necessary to provide any double buffering of the parameter list and thus the memory requirements of the system are further reduced.

The invention is defined with more precision in the appended claims to which reference should now be made.

Preferred embodiments of the invention will now be described in detail by way of example with reference to the accompanying drawings in which:

**Figure 1** shows a schematic diagram of the rendering and texturing system referred to above;

**Figure 2** shows schematically how the display list is ordered in a known tiling system;

**Figure 3** shows how the display list is ordered in accordance with the macro-tiling embodying the invention;

**Figure 4a, b, and c** show alternative arrangements for the allocation of data in blocks in the display list in accordance with an embodiment of the invention;

**Figure 5** shows a further embodiment of the display list;

**Figure 6** shows schematically the z depths for four polygons and z compression by storage of plane parameters for the four polygons;

Figure 7 shows graphically the results of run-length encoding.

Figure 8 shows a block diagram of an example of the type of rendering system with which the present invention may be used;

Figure 9 shows the system of figure 8 in an embodiment of one aspect of the present invention;

Figure 10 is a flow chart illustrating the memory management used by an embodiment of the invention;

Figure 11 is a flow chart for sub-process of the flow chart figure 10; and

Figure 12 is a block diagram of the tile accelerator used in the embodiment of the invention.

In the system of Figure 1, polygons which are to be displayed are broken down into planar triangles. The data representing these is transformed by either hardware or software into a screen space representation and is then stored in local memory. The tiling process then performed by the tile accelerator 2 creates a list of pointers to the transformed triangle data for each tile in the scene. This is then stored in the display list along with parameters pertaining to the object of which the triangle is a portion. This is shown schematically in Figure 2 in which the region headers 10 correspond to the identities of the tiles 1-8 shown in the Figure. As can be seen, the tiles cover two triangles named object A and object B. Object A is visible in tiles 1, 2, 5, 6 and 7 and object B is visible in tiles 3 and 7.

A display list 12 receives from the tile data for the objects visible in each tile and links are provided to the parameters associated with each object, namely parameters A and parameters B in this example. Thus, it can be seen that the object pointers for tiles 1, 2, 5, 6 and 7 point to parameters A and the object pointers for tiles 3 and 7 point to parameters B. When this data is fed to the image

synthesis processor 6 and texturing and shading processor 8, the correct image data will be applied to the pixels of each tile.

5 If the image is complex in relation to the size of the display list it may not be possible to store all the tile data in the display list simultaneously and technique which is referred to as z/frame buffer load and store is used. Therefore, the display list will be loaded with data by the tile accelerator for all the tiles until it is  
10 substantially full. This will only represent a portion of the image data. This data is then rendered by the ISP and TSP and the result stored in local memory. This is called a partial render. The image is rendered one tile at a time in the normal way. After rendering of each tile, the  
15 internal state of the ISP (containing z, stencil and object address information) is written out to external memory and image data from the TSP is written to the frame buffer. These two operations together comprise a z/frame buffer store. At the end of partial render, a render and  
20 a z frame buffer store operation has occurred for each tile on the screen.

At this point, the first part of the display list is discarded (as this data has been used for rendering the tile), and the memory is used to receive further output  
25 from the tile accelerator. When the display list memory fills up again, or when the tile accelerator reaches the end of the scene data, another partial render occurs. As processing begins for each tile in turn, the previously stored ISP data information is reloaded from external  
30 memory and, if necessary, for the blending operations to be performed, the image data from the frame buffer is loaded into the TSP. This is the z/frame buffer load operation. The next portion of the display list can then be rendered in that tile as if it had never been  
35 interrupted. Finally, after the tile is rendered, a z/frame buffer store will save the necessary data

information to allow another cycle of partial renders to occur.

By working in this way, it is possible to render a scene of great complexity whilst using only a limited  
5 amount of memory.

Figure 3 shows how the memory of the display list is allocated in an embodiment of the invention. It can be seen in Figure 3 that the image portion shown in Figure 1 is, in this instance, divided into two macro-tiles headed  
10 macro-tile 1 and macro-tile 2. The display list is divided into two portions, the first portion being for macro-tile 1 and the second portion for macro-tile 2. The region headers for the tiles in macro-tile 1 all feed into the object pointers for macro-tile 1. As only object A is  
15 visible in macro-tile 1, only parameters  $A_1$  need to be stored in this memory block. Thus, it can be seen that the object pointers for tiles 1, 2, 5 and 6 all point into parameters  $A_1$ .

In macro-tile 2 both object A and object B are  
20 visible, both being present in tile 7. Thus, parameters  $A_2$  and parameters  $B_2$  have to be stored in the memory block allocation for macro-tile 2. Thus it can be seen that the object pointer for tile 7 points to parameters  $A_2$  and parameters  $B_2$  whilst the object pointer for tile 3 points  
25 only to parameters  $B_2$ .

The size of macro-tiles can be predetermined or a system can be implemented which defines the size of macro-tiles in dependence on the complexity of the scene. The macro-tiling process will be performed downstream of the  
30 basic tiling operation by the tiling accelerator 2.

The memory available in the display list storage is first divided into allocation blocks ("blocks") and the address of each block is stored in a FIFO or stack known as the free store. The choice of block sizes is a  
35 compromise between space wastage with large blocks and the storage requirements involved in managing a large number

of smaller blocks. It is also influenced by the properties of the memory technology being used.

After this, the screen which has already been divided up into tiles is then divided into macro-tiles. In their  
5 simplest form, all macro-tiles on the screen are the same size which could correspond at its smallest to a single tile. In the example given in Figure 3, a macro-tile size  $2 \times 2$  is used for simplicity of illustration, although in practice the size is likely to be larger.

10 Initially, all allocation blocks are in the free store (not illustrated). The tiling accelerator identifies a set of tiles (and macro-tiles) that a particular object, e.g., a triangle, intersects. Then, for each macro-tile, in turn, when the first object data  
15 that crosses the macro-tile is identified, a block from the free store is claimed and becomes associated with that macro-tile. Object data is then stored in that block. Further blocks are claimed from the free store when there is insufficient space in the current blocks to store a  
20 piece of data.

Once a macro-tile has been rendered by the image synthesis processor 6 and texturing and shading processor  
8, the memory block associated with it can be returned to the free store. In order to do this, blocks need somehow  
25 to be linked to the macro-tile with which they are associated. One possibility to implement this is to use a word in each allocation block to create a linked list of blocks. Alternatively, rather than adding and removing allocation blocks from the free store list, they can be  
30 marked as used or unused *in situ*, with links between used blocks created within the free store structure.

As explained above, the tile accelerator is used to generate a set of tiles and strip mask information for a strip of triangles in the usual way. Additionally, it  
35 determines which macro-tiles the strip falls into. The tiling accelerator 2 then writes parameter data in



independent lists for each macro-tile in which the strip is visible. This parameter data is written to the display list in the appropriate place for each macro-tile.

Clearly, some duplication of data will arise. However,  
5 this allows all the memory associated with a particular macro-tile to be released after that macro-tile has been rendered without affecting the display lists of other macro-tiles.

At the beginning of the first frame of image data to  
10 be displayed, tiling begins. Object pointers for the objects visible in each macro-tile are produced and written to a block from the display list allocated to that macro-tile along with the relevant object parameters for each object visible in that macro-tile. This process  
15 continues until the display list is substantially full or some threshold has been exceeded, e.g., 75% of the allocation blocks have been used. When this triggering level is reached, the systems selects a macro-tile and begins to render it into the frame buffer using the image  
20 synthesis processor and the texture and shading processor. The manner in which the macro-tile is selected is discussed in more detail below. The decision may depend on a number of factors. Typically the macro-tile that currently uses the largest number of allocation blocks  
25 will be selected since this frees up the maximum amount of memory for new data. Tiling continues in parallel with rendering for the remainder of data for the first frame. The 75% threshold is used because this leaves sufficient unused memory to be allocated by the continuing tiling  
30 process whilst other memory is rendered and subsequently returned to the free store. So that tiling can be performed into the macro-tile which is being rendered, its region header (the pointer to the macro-tile's pointer lists and allocation block(s)) is exchanged for a pointer  
35 to an empty tile. Rendering can then continue for the data for the previous macro-tile and the tiling

accelerator can write data into the new allocation block for the empty one. When the rendering of a macro-tile is complete, its allocation blocks are returned to the free store. Then, provided that the memory use is above  
5 threshold value, another macro-tile is chosen and another rendering operation begins. The 75% threshold is an adjustable parameter. This is because the ISP and TSP can perform operations such as deferred texturing more efficiently when a larger part of the display list is in  
10 memory (and most efficiently when the whole list is in memory). The threshold is adjustable to enable a balance to be found between the conflicting requirements of large display list size, and adequate space into which to tile the remainder of the scene.

15 When tiling of a frame of image data is complete, the system can then begin to tile the next frame, even while the rendering of the previous frame is still in progress. The allocation blocks associated with macro-tiles from the new frame must be distinct from those associated with  
20 macro-tiles from the previous frame. This can be achieved by maintaining a second set of blocks independent of those used by the previous frame. This would be achieved by maintaining independently a second set of region headers. Allocation blocks are taken from the same free store for  
25 all frames, which means that it is not necessary to double the size of the display list. This gives a further saving in memory usage.

Thus, it can be seen that the z/frame buffer load and store processes occur at the beginning and end of each  
30 render respectively. They only take place with the tiles in a macro-tile that has been rendered. There is no need for a load operation before the first rendering and no need for a store after the last rendering. It can be seen that the efficiency and memory usage arises from rendering  
35 a macro-tile into the frame buffer when the number of allocation blocks used reaches some predefined threshold

and that upon completion of such a render, the memory blocks associated with the macro-tile rendered are freed up, making them available for further parameter storage.

5 The object parameters which are stored are typically coordinates of triangles where these are the objects being rendered, object colours, etc. As can be seen, both object parameters such as these and object pointers are stored in the allocation block for each macro-tile. This allows parameters and pointers to be discarded easily when  
10 they are no longer required. The proximity of parameter and pointer data in the memory can lead to a significant reduction in the number of page breaks in the memory system if sufficiently small allocation blocks or an appropriate layout with larger blocks is used. Page  
15 breaks reduce the performance of the memory system and are caused when consecutive accesses to the memory hardware are not within the same page. The storage of object pointers and parameter data within the same page wherever possible is an effective way to minimise these page  
20 breaks.

Various schemes from layouts of pointers and parameters within the allocation block are possible. Some examples are shown in Figure 4a, b and c. The simplest schemes allocate pointers in small pointer blocks at the  
25 head of the allocation block to optimise reading and writing. The last pointer in the block is a link to the next pointer block where a linked list is required. In Figure 4a, pointers are allocated from one end of the block and parameters at the other. When the block is full  
30 the two meet in the middle. In Figure 4b, the parameters and pointers are interleaved and the allocation block is filled from beginning to end. A further variation as shown in Figure 4c is more suitable for allocation blocks larger than one page. It attempts to keep object pointers  
35 in the same data block as the parameter data to which they

point but at the expense of some wasted space in the object pointer blocks.

Because allocation blocks are associated with macro-tiles and each macro-tile contains a number of tiles,  
5 there will be multiple linked lists of object pointer blocks within each allocation block.

Where a scene consists of large objects and the macro-tiles each contain a large number of tiles, it may be found that a large amount of space in the allocation  
10 blocks is wasted by pointer blocks that are mostly empty, and whose pointers all point to the same objects.

Figure 5 shows a scheme where a single list is used with a tile mask indicating in which of the tiles of the macro-tile the object is visible. Strip masks are  
15 separate from object pointers and are packed. The step of fetching the image synthesis processor parameters can skip to the relevant strip marks by counting the number of bits set in the tile mask.

In all the allocation schemes described above it is  
20 possible to compress object pointers by representing them using the appropriate number of least significant bits, since the pointer is always pointing to parameter data within the same allocation block. Pointers to objects outside the block (e.g., in a global list) would need to  
25 have the most significant bits of the address supplied as well. This could be by using a larger pointer.

A very large triangle or strip, i.e., one that crosses many macro-tiles can cause a problem for the scheme described above. This is because the data is  
30 replicated in many macro-tiles and thus uses a large amount of memory. One solution to this problem is to create a global list in which large objects are stored. References to these objects are then written into each macro-tile's pointer list in the normal way to ensure that  
35 the order of objects is preserved.

The object could be deemed to be large when it is visible in a number of macro-tiles greater than some threshold value. This threshold could be chosen to achieve an optimal trade-off between parameter replication, memory used by the global list and recycling of parameter memory. As the parameter data in the global list may be referenced in the pointer lists of many macro-tiles it remains in memory for a relatively long period of time. It is therefore important that the global list is not allowed to grow so large as to adversely affect the operation of the macro-tiling. A crude solution would be to render all macro-tiles when the size of the global list exceeds a threshold. This would allow all allocation blocks associated with a global list to be returned to the free store. A more elaborate scheme could employ reference counting, thereby allowing global list allocation blocks to be released as soon as the macro-tiles that reference them have been rendered.

The choice of which macro-tile to render next is generally determined by the maximum amount of memory which will be released back to the free store. This increases efficiency by ensuring that macro-tiles are able to develop a reasonable amount of complexity between renders and also reduces the total number of z/frame buffer load and store operations that are required.

When the display list contains blocks associated with macro-tiles for a number of frames, it is essential that all tiles belonging to one frame are rendered completely before rendering any tile from the next.

It is easy to imagine where a busy macro-tile is being rendered and, at the same time, the tiling accelerator is writing more data into the same macro-tile. This could arise where a small area of the screen contains a particularly detailed object. In such a situation it may make sense to weight the choice of the next macro-tile to render in favour of repeating the same one. This would

allow z load and store to be omitted for at least one tile in the macro-tile or for more than one if a small on-chip cache is present. For the greatest benefit it must be possible to render the tiles in the macro-tile in an arbitrary order.

When the global list becomes uncomfortably large, an extension of the reference counting scheme described above would allow a macro-tile to be chosen for rendering that will release the largest possible number of global list allocation blocks.

An alternative implementation would be to use a hierarchal macro-tile system using nested macro-tiles of progressively increasing size. This aim is to reduce replication of parameter data by selecting an appropriate level of the hierarchy for storage, depending on the number of tiles in which an object appears. Object pointers would be stored at the lowest level only in order to preserve the order of the objects.

The macro-tiled memory management technique described above reduces the bandwidth requirements of the system considerably, but cannot eliminate z/frame buffer load and store entirely. It should also be noted that the stored z information occupies additional memory space. Compression of the z and frame buffer data reduces memory bandwidth and storage requirements still further.

A variety of compression techniques can be used that to reduce the amount of space occupied by the saved data. The efficiencies of different compression methods are affected by the nature of the data being compressed, such that a single scheme is unlikely to offer effective compression in all situations. Therefore, a hierarchical system is proposed, in which the system will fall back to an alternative method of compression if the current scheme is not appropriate to the current data set.

Alternatively, any one of the proposed methods might be

used in isolation, at the risk of offering poor data compression (or even expansion) in some cases.

**1) Storage of face plane equations of surfaces for z compression**

5       The z values in a tile are comprised of the depths of the triangles that are visible at each pixel. Although there may be many triangles processed for each tile, it is often the case that only a small number of triangles are visible, and therefore contributing to the depth values.

10      In the case of the z-clear polygons that occur at the beginning of most images, one polygon will define the initial z state for the entire tile.

      For a given polygon, the z information is defined by the three plane parameters, A, B, and C, calculated by the

15      ISP. The parameters are used in the equation  $\text{depth} = Ax + By + C$ , where x and y are the position of the pixel in question, with (0,0) at the top left corner of the tile.

      Storage of the A, B and C parameters requires 12

20      bytes of data per triangle. An additional bit map value is stored per-pixel to select between different sets of parameters. An example where the z depths are defined by four polygons is shown in Figure 6.

      Compressing the data requires knowing how many

25      polygons contribute to the z values, and the ability to recall the plane parameters for those polygons. This data could be computed from the stored object identity tags at the point where the data is output to the TSP. In the case where more than eight polygons contribute to the

30      depths, the system would fall back to one of the alternative schemes mentioned below.

      Compression rates for different numbers of triangles contributing to the depth information in a rectangular tile of 32 x 16 pixels are as follows:

Num Polys	ISP Data Size	Bit map data size	Total data size
1	12 bytes	0	12 bytes
2	24 bytes	64 bytes	88 bytes
3-4	48 bytes	128 bytes	176 bytes
5-8	96 bytes	256 bytes	352 bytes

The raw floating point z data would require 2048 bytes. The method can therefore be seen to give a considerable compression ratio of between 5 and 128 times.

Reloading these depth values would require calculations similar to those that the ISP performs when processing a triangle. The depth values computed in this way would be written into the depth buffer according to the bit map data in order to reconstruct the original set of values.

The example above allows the Z depth to be defined by between one and eight triangles, although other ranges could be used. If the number of triangles falls outside this range, the system can fall back to an alternative system, described below.

## 2) Run length encoding for Z compression

Run length encoding is commonly used when data contains large numbers of consecutive items with the same value. Numbers are stored in pairs, so that, for example, the encoding 8:0.5, 24:0.8 would indicate a row of pixels where the first eight had a depth value of 0.5, and the following 24, presumably belonging to a different triangle, have a depth value of 0.8.

An alternative type of run length encoding will usually prove more effective for encoding depth information, where it is quite likely that triangles are not "flat on" and therefore have depth values that change uniformly across the tile. This situation is shown in Figure 7. Since triangles are planar, it is found that



the "delta" values (the differences in depth between one pixel and the next) are constant. In fact, when moving across the tile in the x direction, the delta value is equal to the ISP's "A" parameter for that polygon. Run length encoding the x delta values along lines of the tile (constant y) gives a compression ratio of approximately 2 times with typical data (from experimental results).

Run length encoding (of values rather than deltas) is also highly appropriate for stencil and "tag" (the identity of the object at a particular pixel position) values, which may need to be stored along with the depth information to fully record the state of the ISP.

A disadvantage of run length encoding is that it is possible to have a situation where the data size actually expands. This would occur, for example, where no adjacent pixels are identical. The compressed data could actually double in size, since it would be necessary to store a counter indicating that each value occurs only once. In this case, the Z compression system would fall back to storing uncompressed z values.

### **3) Frame buffer compression**

The frame buffer memory is read directly by the hardware that produces the video image, and therefore the final output cannot be compressed. However, if data is stored and loaded a number of times in the process of composing the image, memory bandwidth can be conserved by compressing the intermediate data. Standard image compression techniques, including Run Length Encoding, J-PEG, Wavelet, or Vector Quantisation (VQ) may be used.

### **4) Write Z only when changed, read TSP only on demand**

A final optimisation may make it possible to avoid some z stores and frame buffer loads altogether.

A "dirty" flag indicates that data has been modified since it was last read. In the case of z buffer

load/store, the dirty flag would be cleared when a z load occurs, and set when the z data is modified by the ISP. When a z store is requested, it is only necessary to actually save the data if the dirty bit is set.

5           A similar optimisation can be applied to frame buffer loads and stores. It is only necessary to load pixels from the frame buffer when they are required by a particular blending operation, and this situation may not always arise. Similarly, a dirty flag would indicate  
10          pixels or groups of pixels which have been modified, and only those pixels would be transferred by a frame buffer store operation.

**5)    Allowing the application to access the "z-buffer" if it wants to**

15          Some rendering systems such as the Imagination Technologies Power Vr System do not normally need to maintain a z buffer for the entire frame. This saves memory, and bandwidth. Some applications, however, expect to see a z buffer, and attempt to manipulate it themselves  
20          in order to achieve special graphical effects. Although the data stored by the z buffer store operations is not a z buffer in the conventional sense, it is possible to make it appear as such to the application, by taking care over the layout of data in memory. If necessary additional z  
25          buffer store operations can be performed to ensure that the saved z values reflect the state of the imaginary z buffer at the appropriate times.

          A conventional z buffer contains raw, uncompressed z data. Clearly, any application that attempted to read a z  
30          buffer compressed using the methods described above would be unable to interpret it correctly. It is essential therefore, that there is a facility to either switch off the compression for those applications that require it, or to provide additional hardware that presents a "virtual" z

buffer to the application, and performs compression and decompression transparently.

Diagrams illustrating an embodiment of the invention now follow. Figure 8 shows the block diagram of the type of system to which the invention is applied. This comprises a slightly amplified version of figure 1 and shows a tile accelerator 2 performing tiling operations which are supplied to a display list memory 4. This in turn supplies data to an image synthesis processor 6 for performing ray/surface intersection tests. After this data is supplied to a texture and shading processor 8 which applies texture and shading data, from a texture in memory 9 before storing the textured image data in frame buffer memory 11. This apparatus operates in accordance with the description given earlier with relation to figure 1.

Figure 9 shows this system figure 8 modified by circuitry to perform a Z/frame buffer load and store operation. In this, a Z buffer memory 16 is linked to the image synthesis processor 6 via a z compression/decompression unit 18. This comes into operation when the system is rendering a complex scene and the display list memory is not large enough to contain all the surfaces which need to be processed for a particular tile. This process has been described previously. However, when operating using the system of figure 9, the display list will be loaded with data by the tile accelerator for all the tiles until it is substantially full. This may, however, only represent a portion of the initial data. The image is rendered one tile at a time by the image synthesis processor 6. The output data for each tile is provided to the texture and shading processor 8 which uses data from the texture memory 9 to texture the tile and supply it to the frame buffer memory 11.

At the same time, because the image data was incomplete, the result from the image synthesis processor 6 is stored to z buffer memory 16 via the compression/decompression unit 18 for temporary storage. The rendering of the remaining tiles then continues with the incomplete image data until all the tiles have been rendered and stored in frame buffer memory and in the z buffer memory 16.

The first part of the display list is then discarded and the additional image data read into it. As processing is performed for each tile in turn by the image synthesis processor 6 the relevant portion of data from the z buffer memory 16 is loaded via the z compression/decompression unit 18 so that it can be combined with the new image data from the display list memory 4. The new data for each tile in turn is then fed to the texture and shading processor 8 which combines it with data from texture memory 9 before supplying it to the frame buffer 11.

This process continues for all the tiles in the scene and until all the image data has been rendered.

Thus, it can be seen that the z buffer memory fills a temporary store which enables a smaller display list memory to be used than would be necessary for rendering particularly complex scenes. The compression/decompression unit 18 is optional but its use enables a smaller z buffer memory to be employed.

A frame buffer compression/decompression unit 20 is provided as an additional link between the texture and shading processor and the frame buffer memory. Preferably, this is used for temporary storage of incomplete image data in the frame buffer memory and for feeding back that data when additional data is provided from the image synthesis processor 6 for a tile which has to be combined with the data already

textured and shaded for that tile. This compression/decompression unit is of course optional.

In figure 10 the memory management system used by the tiling accelerator is illustrated. This should be read in combination with the description of figure 3.

5 At step 30, the region headers are initialised and step 32 the system waits for object data to be input for the regions. At 34 tiles and macro tiles are allocated as illustrated in figure 3 and at 36 its determination is made as to whether or not any new memory blocks are

10 required. If additional memory blocks are not required then object data in macro tiles is stored in memory blocks as previously allocated by the system whilst at 40 the determination is made as to whether or not the last object in the frame has been reached in the

15 allocation of object data to macro tiles and memory. At 42 determination is made as whether or not any unrendered macro tiles exist for the frame if they do then at Step 44 a macro tile is selected for rendering

20 and the rendering commenced at 46 before a determination is made at 48 as to whether or not additional objects are required to be stored and rendered. If the answer is no then at 42 a

determination is made as to whether another tile is to be rendered and if no further macro tiles exist, a

25 determination is made at 50 as to whether another frame or data is to be entered. If it is, the system returns to the top of the flow diagram. If not this process terminates. If the determination at 42 is yes then

30 control flows to fig 6 for new blocks from memory to be loaded.

If at step 36 a determination is made that additional memory is required then at step 52 a determination is made as to whether or not the memory

35 fill threshold has already been exceeded. If it is not

then at 54 new memory blocks are claimed from the free store at 54 and are used at 38 to store object data in macro tiles before proceeding to step 40 as before. If the memory fill threshold has been exceeded then the process skips to step 44 where macro tiles are rendered. If this route is taken at step 48 the determination will be that the system is waiting to store an object. In which case the process will activate step 56 which will wait until memory blocks are available before they can be claimed at step 54 for storing additional objects in macro tiles at step 38.

Figure 11 illustrates how memory blocks are freed and returned to store for future allocation by the tile accelerator. At 80, the system is activated to wait for a signal that a partial render has ben completed. When this is received, memory blocks are returned to store at 82 and the system waits for the end of the next partial render.

Because tiling and rendering occur in parallel, this freeing up of memory blocks is illustrated separately in figure 10, at 46 a render is commenced by sending a signal to the ISP. It doesn't wait for it to complete, because the tile accelerator can be performing more tiling work in that time. This is why there is a threshold (e.g. 79%) at which rendering of macro tiles commences. This leaves the unfilled portion of memory for the tile accelerator to work with.

The system of figure 11 waits for the ISP to indicate that a partial render is complete because it doesn't know the stage in the cycle at which the tile accelerator is and doesn't wish to interrupt it. Figure 11 simply releases what memory it can, such that when steps 56 and 54 in figure 10 are reached. There is a good chance that the memory will be immediately available for use. It is necessary sometimes to wait at 56 because memory will not always be available. This may occur with complex images

and it is then necessary to wait for memory to become available.

Figure 12 shows the ISP 6 and display list memory 4 at figures 8 and 9 but with the tile accelerator 2 shown in more detail. Input data 90 passes to a tiling unit 92 and this data then passes to a macro-tiling unit 94. Output data 95 from this passes to the display list memory 4 and also to a display list memory manager 96. The memory manager 96 is in communication with the free store list 98 and with the display list 4. It receives a signal from the display list 45 to indicate that the memory is e.g. 75% full and in response to this sends a start render signal to the ISP6. When this is complete an end render signal is sent to the memory manager 96 which returns the memory now made available to the free store list 98 for use by further macro tiles. The memory manager continues to receive macro-tile data, whilst previous macro-tiles are being rendered, and allocates memory from the free store list to the objects in these as required.